

# CORE: Augmenting Regenerating-Coding-Based Recovery for Single and Concurrent Failures in Distributed Storage Systems

Runhui Li, Jian Lin, Patrick P. C. Lee

Department of Computer Science and Engineering, The Chinese University of Hong Kong

{rhli, jlin, pclee}@cse.cuhk.edu.hk

**Abstract**—Data availability is critical in distributed storage systems, especially when node failures are prevalent in real life. A key requirement is to minimize the amount of data transferred among nodes when recovering the lost or unavailable data of failed nodes. This paper explores recovery solutions based on regenerating codes, which are shown to provide fault-tolerant storage and minimum recovery bandwidth. Existing optimal regenerating codes are designed for single node failures. We build a system called CORE, which augments existing optimal regenerating codes to support a general number of failures including single and concurrent failures. We theoretically show that CORE achieves the minimum possible recovery bandwidth for most cases. We implement CORE and evaluate our prototype atop a Hadoop HDFS cluster testbed with up to 20 storage nodes. We demonstrate that our CORE prototype conforms to our theoretical findings and achieves recovery bandwidth saving when compared to the conventional recovery approach based on erasure codes.

**Keywords**—regenerating codes, failure recovery, distributed storage systems, coding theory, experiments and implementation

## I. INTRODUCTION

To provide high storage capacity, large-scale distributed storage systems have been widely deployed in enterprises, such as Google File System [14], Amazon Dynamo [10], and Microsoft Azure [4]. In such systems, data is striped across multiple nodes (or servers) that offer local storage space. Nodes are interconnected over a networked environment, in the form of either clustered or wide-area settings.

Ensuring data availability in distributed storage systems is critical, given that node failures are prevalent [14]. Data availability can be achieved via erasure codes (e.g., Reed-Solomon codes [34]), which encode original data and stripe encoded data across multiple nodes. Erasure codes are defined by parameters  $(n, k)$  (where  $k < n$ ), such that if any subset of  $n - k$  out of  $n$  nodes fails, the original data remains accessible by decoding the encoded data stored in other  $k$  surviving nodes. Erasure codes can tolerate multiple failures, while incurring less storage overhead than replication.

In addition to tolerating failures, another crucial availability requirement is to *recover* any lost or unavailable data of failed nodes. Recovery is performed in two scenarios: (i) when the failed nodes are crashed and the permanently lost data need to be restored on new nodes, and (ii) when the unavailable data needs to be accessed by clients before the failures are restored. High-performance recovery is necessary in both scenarios. The conventional recovery approach, which applies to *any*

erasure codes, first reconstructs all original data to obtain the lost/unavailable data. Since the lost/unavailable data usually accounts for only a fraction of original data, previous studies explore how to optimize the recovery performance by minimizing the amount of data involved. One class of approaches is to minimize I/Os, or the amount of data read from disks, based on erasure codes (e.g., [22], [26], [35], [44], [47], [50], [51]). Another class of approaches is to minimize bandwidth, or the amount of data transfer over a network during recovery, based on *regenerating codes* (e.g., [11], [32], [42]), in which each surviving node encodes its stored data and sends encoded data for recovery using network coding [1]. In the scenario where network capacity is limited, minimizing the recovery bandwidth can improve the overall recovery performance. In this work, we focus on exploring the feasibility of deploying regenerating codes in practical distributed storage systems.

However, most existing recovery approaches, including those for minimizing I/Os and bandwidth, are restricted to *single failure* recovery. Although single failures are common, node failures are often correlated and co-occurring in practice, as reported in both clustered storage (e.g., [13], [36]) and wide-area storage (e.g., [5], [16], [29]). To provide tolerance against *concurrent* (multiple) failures, data is usually protected with a high degree of redundancy. For example, Cleversafe [6], a commercial wide-area storage system, use (16,10) erasure codes (i.e., up to 6 out of 16 concurrent failures are tolerable) [31]. Some wide-area storage systems such as OceanStore [27] and CFS [8] employ erasure codes with the even higher double redundancy  $(n, n/2)$ . We believe that in addition to providing fault tolerance, minimizing the recovery bandwidth for concurrent failures will provide additional benefits for today's large-scale distributed storage systems. In addition, concurrent recovery is beneficial to delaying immediate recoveries [2]. That is, we can perform recovery only when the number of failures exceeds a tolerable limit. This avoids unnecessary recoveries should a failure be transient and the data be available shortly (e.g., after rebooting a failed node). Given the importance of concurrent recovery, we thus pose the following questions: (1) Can we achieve bandwidth saving, based on regenerating codes, in recovering a general number of failures including single and concurrent failures? (2) If we can enable regenerating codes to recover concurrent failures, can we seamlessly integrate the solution into a practical distributed storage system?

In this paper, we propose a complete system called CORE,

which supports both single and concurrent failure recovery and aims to minimize the bandwidth of recovering a *general* number of failures. CORE augments existing optimal regenerating codes (e.g., [32], [42]), which are designed for single failure recovery, to also support concurrent failure recovery. A key feature of CORE is that it retains existing optimal regenerating code constructions and the underlying regenerating-coded data. That is, instead of proposing new code constructions, CORE adds a new recovery scheme atop existing regenerating codes. Our idea is to treat all but one failed nodes as logical surviving nodes. CORE first reconstructs the “virtual” data to be generated by those logical surviving nodes. By combining the virtual data with the real data being generated by the actual surviving nodes, CORE then reconstructs the remaining failed node using existing optimal regenerating codes. We apply the same idea for all failed nodes.

In summary, the contributions of this paper are three-fold.

- **Theoretical analysis.** We theoretically show that CORE achieves the minimum bandwidth for a majority of concurrent failure patterns. We also propose extensions to CORE to achieve sub-optimal bandwidth saving even for the remaining concurrent failure patterns. Our analytical study validates that CORE can recover concurrent failure patterns with significant bandwidth saving over conventional recovery based on erasure codes. For example, for (20,10), the bandwidth savings are 36-64% and 25-49% in the optimal and sub-optimal cases, respectively.
- **Implementation.** We implement a prototype of CORE and demonstrate the feasibility of deploying CORE in a practical distributed storage system. To make a case, we choose the Hadoop Distributed File System (HDFS) [40] as a starting point. CORE sits as a layer atop HDFS and supports recovery for a general number of failures. We build CORE atop HDFS by modifying the source code of HDFS and its erasure coding extensions HDFS-RAID [18]. We also adopt a pipelined implementation that parallelizes and speeds up the recovery process.
- **Experiments.** We experiment CORE on an HDFS testbed with up to 20 storage nodes. Our experiments take into account a combination of different factors including network bandwidth, disk I/Os, encoding/decoding overhead. We justify that minimizing bandwidth in recovery plays a key role in improving the overall recovery performance. We show that compared to erasure codes, CORE achieves recovery throughput gains with up to  $3.4\times$  for single failures and up to  $2.3\times$  for concurrent failures. Our experimental results conform to our theoretical findings. Our prototype also maintains the performance of striping replicas into encoded data, an operation that is included in original HDFS-RAID, when regenerating codes are used.

The rest of the paper proceeds as follows. Section II first formulates our system model. Section III motivates how CORE reduces bandwidth of conventional recovery. Section IV describes the design of CORE and presents our theoretical and analysis findings. Section V describes the implementation

TABLE I  
MAJOR NOTATION USED IN THIS PAPER.

$n$	number of nodes
$N_i$	the $i$ -th node ( $0 \leq i \leq n-1$ )
$k$	number of data nodes
$r$	number of symbols per strip
$t$	number of concurrent failures ( $1 \leq t \leq n-k$ )
$M$	size of original data stored in a stripe
$s_{i,j}$	the $j$ -th stored symbol in a stripe of node $N_i$ ( $0 \leq i \leq n-1, 0 \leq j \leq r$ )
$e_{i,i'}$	encoded symbol from surviving node $N_i$ used to recover lost data of failed node $N_{i'}$ ( $0 \leq i, i' \leq n-1$ )

details of CORE. Section VI presents experimental results. Section VII reviews related work, and Section VIII concludes this paper and presents future work.

## II. SYSTEM MODEL

We formulate the recovery problem in a distributed storage system. We also provide an overview of regenerating codes, and show how they can improve the recovery performance.

### A. Basics

We first define the terminologies and notation. Table I summarizes the major notation used in this paper. We consider a distributed storage system composed of a collection of *nodes*, each of which refers to a physical storage device. The storage system contains  $n$  nodes labeled by  $N_0, N_1, \dots, N_{n-1}$ , in which  $k$  nodes (called *data nodes*) store the original (uncoded) data and the remaining  $n-k$  nodes (called *parity nodes*) store parity (coded) data. The coding structure is *systematic*, meaning that the original data is kept in storage.

Figure 1 shows an example of a distributed storage system, which is also consistent with the erasure-coded design of HDFS-RAID [18]. Each node stores a number of *blocks*. A block is the basic unit of read/write operations in a storage system. It is called a data block if it holds original data, or a parity block if it holds parity data. To store data/parity information, each block is partitioned into fixed-size *strips*, each of which contains  $r$  *symbols*. A symbol is the basic unit of encoding/decoding operations. A *stripe* is a collection of strips on  $k$  data nodes and the corresponding encoded strips on  $n-k$  parity nodes. A data (parity) block contains all strips of data (parity) symbols. For load balancing reasons the identities of the data/parity nodes are rotated so that the data and parity blocks are evenly distributed across nodes [26], [31].

Each stripe is independently encoded. Our discussion thus focuses on a single stripe and our recovery scheme will operate on a per-stripe basis. Let  $M$  be the total amount of original uncoded data stored in a stripe. Let  $s_{i,j}$  be a stored symbol of node  $N_i$  at offset  $j$  in a stripe, where  $i = 0, 1, \dots, n-1$  and  $j = 0, 1, \dots, r-1$ . Each stripe contains  $nr$  stored symbols, which can be formed by multiplying an  $nr \times kr$  *generator matrix* by a vector of  $kr$  original data symbols based on the *Galois field* arithmetic, whose implementation details can be found in the prior study [15]. In this work, we focus on the arithmetic operations over the Galois field  $\text{GF}(2^8)$ . Note that




Fig. 1. Example of a distributed storage system, where  $n = 6$ ,  $k = 3$ , and  $r = 3$ . We assume that nodes  $N_0$ ,  $N_1$ , and  $N_2$  are data nodes, while  $N_3$ ,  $N_4$ , and  $N_5$  are parity nodes. For load balancing, the identities of data and parity nodes are rotated across different blocks.

our recovery scheme applies to the failures of both data and parity nodes. It treats each stored symbol  $s_{i,j}$  the same way regardless of whether it is a data or parity symbol.

For data availability, we have the storage system employ an  $(n, k)$  code that is *maximum distance separable (MDS)*, meaning that the stored data of any  $k$  out of the  $n$  nodes can be used to reconstruct the original data. That is, an  $(n, k)$  MDS-coded storage system can tolerate any  $n - k$  out of  $n$  concurrent failures. MDS codes also ensure optimal storage efficiency, such that each node stores  $\frac{M}{k}$  units of data per stripe. Reed-Solomon (RS) codes [34] are a classical example of MDS codes. RS codes can be implemented with stripe size  $r = 1$  to minimize the generator matrix size.

### B. Recovery

Our recovery addresses two types of node failures. The first type is the recovery from permanent failures (e.g., due to crashes) where data is permanently lost. In this case, we reconstruct the lost data of the failed nodes on new nodes to minimize the window of vulnerability. Another type is degraded reads to the temporarily unavailable data during transient failures (e.g., due to system reboots or upgrades) or before the permanent failures are restored. The reads are degraded as the unavailable data needs to be reconstructed from the available data of other surviving nodes. In our discussion, we use “lost data” to refer to both permanently lost data and temporarily unavailable data.

We consider the scenario where the storage system activates recovery of lost data when there are a number  $t \geq 1$  of failed nodes. Clearly, we require  $t \leq n - k$ , or the original data will be unrecoverable. We call the set of  $t$  failed nodes the *failure pattern*. The lost data will be reconstructed by the data stored in other surviving nodes.

Our recovery builds on the *relayer* model, in which a relayer daemon coordinates the recovery operation. Figure 2 depicts the relayer model. During recovery, each surviving node performs two steps: (i) *I/O*: it reads its stored data, and (ii) *encode* (for regenerating codes only): it combines the stored data into some linear combinations. The relayer daemon performs three steps: (i) *download*: it downloads the data from some other surviving nodes, (ii) *reconstruction*: it reconstructs the lost data, and (iii) *upload*: it uploads the reconstructed data to the new nodes (for recovery from permanent failures) or

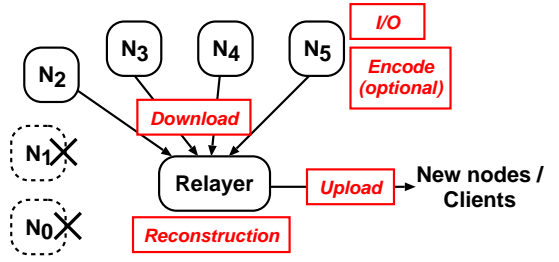


Fig. 2. Recovering nodes  $N_0$  and  $N_1$  using the relayer model.

to the client who requests the data (for degraded reads). We assume that the relayer is reliable during the recovery process.

We argue that the relayer model can be easily fit into practical distributed storage systems. In the case of recovering permanent failures, we can deploy the relayer daemon in different ways, such as in one of the new storage nodes that reconstructs all lost data, in every storage node that reconstructs a subset of lost data, or in separate servers that run outside the storage system. In the case of degraded reads, we can deploy the relayer daemon in each storage client. We note that this relayer model is also used in prior studies in the contexts of peer-to-peer storage [2], data center storage [22], and proxy-based cloud storage [20]. In Section V, we elaborate how the relayer model can be integrated into a distributed storage system.

To improve the recovery performance of a distributed storage system with limited network bandwidth, it is important to minimize the amount of data transferred over the network. If the number of failed nodes is small, the amount of data being downloaded from the surviving nodes is larger than the amount of reconstructed data being uploaded to new nodes (or clients). If we pipeline the download and upload steps (see Section V-B), then the download step becomes the bottleneck. Thus, we focus on optimizing the download step in recovery. Formally, we define the *recovery bandwidth* as the total amount of data being downloaded per stripe from the surviving nodes to the relayer during recovery. Our goal is to minimize the recovery bandwidth.

### C. Regenerating Codes

When an erasure-coded system sees failures, *conventional recovery* is used, meaning that the relayer downloads data from any  $k$  surviving nodes to first reconstruct all original data and then return the lost data. The amount of data being downloaded is equal to the amount of original data being stored (i.e.,  $M$  per stripe). Note that some proposals allow less data to be read for some erasure codes under specific conditions (see Section VII). However, conventional recovery applies to *any MDS erasure code and any number of failures no more than  $n - k$* . In this paper, when we refer to erasure codes, we assume that conventional recovery is used.

We consider a special class of codes called *regenerating codes* [11] that enables the relayer to transfer less than the amount of original data being stored. Regenerating codes build on network coding [1], in which during recovery, surviving



nodes send encoded symbols that are computed by the linear combinations of their stored symbols, and then the encoded symbols are used to reconstruct the lost data. It is shown that regenerating codes lie on an optimal tradeoff curve between storage cost and recovery bandwidth [11]. There are two extreme points: *minimum storage regenerating (MSR)* codes, in which each node stores the minimum amount of data on the tradeoff curve, and *minimum bandwidth regenerating (MBR)* codes, in which the bandwidth is minimized. Note that MSR codes have the same optimal storage efficiency as MDS erasure codes such as RS codes, while MBR codes minimize bandwidth at the expense of higher storage overhead. In this work, we focus on MSR codes.

Existing optimal MSR codes are designed for recovering a single failure, as described below. First, the strip size has  $r = n - k$  symbols to achieve the minimum possible bandwidth. During recovery, the relay downloads one *encoded* symbol from each of the  $n - 1$  surviving nodes<sup>1</sup>. Let  $e_{i,i'}$  be the encoded symbol downloaded from node  $N_i$  and used to reconstruct data for the failed node  $N_{i'}$ . Each encoded symbol  $e_{i,i'}$  is a function of the symbols  $s_{i,0}, s_{i,1}, \dots, s_{i,r-1}$  stored in the surviving node  $N_i$ , and has the same size as each stored symbol. Using the encoded symbols, the relay reconstructs the lost symbols of the failed node  $N_{i'}$ . MSR codes achieve the minimum recovery bandwidth (denoted by  $\gamma_{MSR}$ ) for single failure recovery given by [11]:

$$\gamma_{MSR} = \frac{M(n-1)}{k(n-k)}. \quad (1)$$

However, existing studies on regenerating codes are limited in different aspects, which we further discuss in Section VII. To summarize, most recovery approaches focus on single failures. If more than one node fails, the optimal MSR code constructions cannot achieve the saving shown in Equation (1) by connecting to  $n - 1$  surviving nodes. To recover concurrent failures, a straightforward approach is to resort to conventional recovery and download the size of original data from any  $k$  surviving nodes. This paper explores if we can achieve recovery bandwidth saving for concurrent failures as well.

### III. MOTIVATING EXAMPLE

Before we describe the design of CORE, we first motivate via an example how CORE reduces the recovery bandwidth over conventional recovery for concurrent failures. The design details of CORE will be refined in Section IV.

We consider an MDS code with  $n = 6$  and  $k = 3$ . Suppose that we store a data object of size  $M$  that corresponds to a stripe of original data symbols. For erasure codes, the strip size is  $r = 1$  symbol, and the symbol size is  $\frac{M}{3}$ . For regenerating codes, the strip size is set to  $r = n - k = 3$  symbols, and hence the symbol size is  $\frac{M}{9}$ . Suppose now both nodes  $N_0$  and  $N_1$  fail. Our goal is to reconstruct their lost data.

<sup>1</sup>There are MSR code constructions (e.g., [32], [42]) that can download encoded symbols from less than  $n - 1$  surviving nodes at the expense of higher recovery bandwidth. In this work, we only focus on the case where  $n - 1$  surviving nodes are connected.

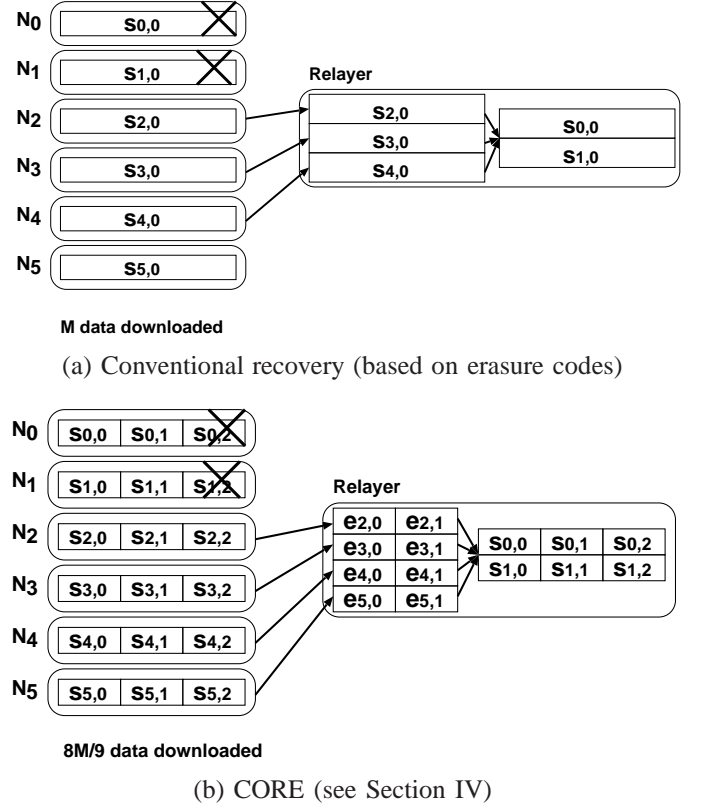


Fig. 3. Comparisons of conventional recovery and CORE.

We first consider conventional recovery based on erasure codes, whose core idea is to first reconstruct all original data. Thus, the relay downloads the size of original data from any  $k = 3$  nodes (e.g.,  $N_2, N_3, N_4$ ). Figure 3(a) shows the conventional recovery, in which the relay reconstructs all three original symbols to regenerate the data for the two failed nodes simultaneously. Thus, the total amount of data downloaded is  $M$ .

We now consider how CORE applies concurrent recovery. Here, we consider the baseline approach of CORE (see Section IV-A). Figure 3(b) shows the main idea, in which the relay now downloads two encoded symbols  $e_{i,0}$  and  $e_{i,1}$  from each of the four surviving nodes  $N_i$  ( $i = 2, 3, 4, 5$ ), such that CORE form a system of equations in  $e_{i,0}$ 's and  $e_{i,1}$ 's to reconstruct the lost data of nodes  $N_0$  and  $N_1$ . The total amount of data downloaded is  $\frac{8M}{9}$ , which is one symbol size less than that of conventional recovery. We point out that the bandwidth saving of CORE can be even higher for some parameters. In general, when the baseline approach of CORE applies concurrent recovery for  $t$  failures, the relay downloads  $t$  encoded symbols from each of the  $n - t$  surviving nodes. We elaborate the details in the next section.

### IV. DESIGN OF CORE

CORE builds on existing MSR code constructions that are designed for single failure recovery with parameters  $(n, k)$ . CORE has two major design goals. First, CORE preserves existing code constructions and stored data. That is, we still

have data be striped and stored with existing MSR code constructions, while CORE sits as a layer atop existing MSR code constructions and enables efficient recovery for both single and concurrent failures. The optimal storage efficiency of MSR codes is still preserved. Second, CORE aims to minimize recovery bandwidth for a variable number  $t \leq n - k$  of concurrent failures, without requiring  $t$  to be fixed before a code is constructed and the data is stored.

In this section, we first describe the baseline approach of CORE, in which we extend the existing optimal solution of single failure recovery to support concurrent failure recovery (Section IV-A). We note that the baseline approach of CORE is not applicable in a small proportion of failure patterns, so we propose a simple extension that still provides bandwidth reduction for such cases (Section IV-B). We present theoretical results that show that CORE can reach the optimal point in a majority of failure patterns (Section IV-C). Finally, we analyze the recovery bandwidth saving of CORE (Section IV-D).

#### A. Baseline Approach of CORE

We first provide the background of existing MSR code constructions on which CORE is developed. We then define the building blocks of CORE, and explain how CORE uses these building blocks to support concurrent failure recovery.

**Background.** CORE can build on existing optimal MSR code constructions including Interference Alignment (IA) codes [42] and Product-Matrix (PM) codes [32]. We here provide a high-level overview of how IA codes work, while PM codes have a similar idea. IA codes extend the idea of aligning interference signals in wireless communication into failure recovery in distributed storage systems. Recall that each stripe in regenerating codes contains  $k(n - k)$  original data symbols (see Section II-A). Each stored symbol is a linear combination of the  $k(n - k)$  original data symbols. Suppose that a data node fails (the similar idea also applies for parity nodes). The  $n - 1$  surviving nodes compute the  $n - 1$  encoded symbols (denoted by  $\mathbf{y} = (y_1, \dots, y_{n-1})^T$ ). The relay downloads the  $n - 1$  encoded symbols and reconstructs the  $n - k$  lost data symbols (denoted by  $\mathbf{x}_1 = (x_1, \dots, x_{n-k})^T$ ) of the failed node. There are other  $(k-1)(n-k)$  data symbols (denoted by  $\mathbf{x}_2 = (x_{(n-k)+1}, \dots, x_{k(n-k)})^T$ ) that do not need to be regenerated and can be viewed as interference signals. We can express  $\mathbf{y}$  as a system of equations in  $\mathbf{x}_1$  and  $\mathbf{x}_2$  as:

$$(\mathbf{A} \mid \mathbf{B}) \begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{pmatrix} = \mathbf{y},$$

for some coefficient matrices  $\mathbf{A}$  and  $\mathbf{B}$  of sizes  $(n-1) \times (n-k)$  and  $(n-1) \times (k-1)(n-k)$ , respectively. By elementary row operations, we can transform the system of equations into:

$$\begin{pmatrix} \mathbf{A}' & \mathbf{0} \\ \mathbf{0} & \mathbf{B}' \end{pmatrix} \begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{pmatrix} = \mathbf{y}',$$

for transformed vector  $\mathbf{y}'$  and transformed matrices  $\mathbf{A}'$  and  $\mathbf{B}'$  of sizes  $(n-k) \times (n-k)$  and  $(k-1)(n-k) \times (k-1)(n-k)$ , respectively. Note that IA codes ensure that there exists some




Fig. 4. An example of how the relay downloads real and virtual symbols for a (6,3) code when there are two failed nodes  $N_0$  and  $N_1$ . Here,  $e_{1,0}$  and  $e_{0,1}$  are the virtual symbols.

transformation that makes  $\mathbf{A}'$  an invertible matrix, so that  $\mathbf{x}_1$  (i.e., the lost symbols) can be uniquely solved.

IA codes design the generator matrix that satisfies the above properties. PM codes have a similar idea using a different generator matrix design. We refer readers to [32], [42] for their mathematical details on the generator matrix design.

Note that both IA and PM codes have parameter constraints. IA codes require  $n \geq 2k$ , and PM codes require  $n \geq 2k - 1$ . In this work, we mainly focus on the double redundancy  $n = 2k$ , which is also considered in state-of-the-art distributed storage systems such as OceanStore [27] and CFS [8]. While the redundancy overhead is higher than traditional RAID-5 and RAID-6 codes for large  $(n, k)$ , it remains less than traditional 3-way replication used in production storage systems such as GFS [14] and HDFS [40].

**Building blocks.** Our observation is that any optimal MSR code construction can be defined by two functions. Let  $\text{Enc}_{i,j}$  be the encoding function that is called by node  $N_i$  to generate an encoded symbol for the failed node  $N_j$  using the  $r = n - k$  stored symbols in node  $N_i$  as inputs; let  $\text{Rec}_i$  be the reconstruction function that returns the set of  $n - k$  stored symbols of a failed node  $N_i$  using the encoded symbols from the other  $n - 1$  surviving nodes as inputs. Both  $\text{Enc}$  and  $\text{Rec}$  define the operations of linear combinations of the stored symbols  $s_{i,j}$ 's, depending on the specific code construction. From the above discussion,  $\text{Enc}$  is to construct the encoded symbols  $\mathbf{y}$ , while  $\text{Rec}$  is to reconstruct the lost symbols  $\mathbf{x}_1$ .

CORE works for *any* construction of optimal MSR codes, as long as the functions  $\text{Enc}$  and  $\text{Rec}$  are well-defined. The two functions  $\text{Enc}$  and  $\text{Rec}$  form the building blocks of CORE.

**Main idea of the baseline approach.** We consider two types of encoded symbols to be downloaded for recovery: *real symbols* and *virtual symbols*. To recover each of the  $t$  failed nodes, the relay still operates as if it connects to  $n - 1$  nodes, but this time it represents the symbols to be downloaded from the failed nodes as virtual symbols, while still downloading the symbols from the remaining  $n - t$  surviving nodes as real symbols. Now, using  $\text{Enc}$  and  $\text{Rec}$ , we reconstruct each virtual symbol as a function of the downloaded real symbols. Finally, using the downloaded real symbols and the reconstructed virtual symbols, we can reconstruct the lost stored symbols in the failed nodes.

**Example.** We depict our idea using Figure 4, which shows a (6,3) code and has failures  $N_0$  and  $N_1$ . The two encoded symbols  $e_{1,0}$  and  $e_{0,1}$  are virtual symbols, and the rest are real symbols. We can express  $e_{1,0}$  and  $e_{0,1}$  based on the functions Enc and Rec for single failure recovery as:

$$\begin{aligned} e_{1,0} &= \text{Enc}_{1,0}(s_{1,0}, s_{1,1}, s_{1,2}) \\ &= \text{Enc}_{1,0}(\text{Rec}_1(e_{0,1}, e_{2,1}, e_{3,1}, e_{4,1}, e_{5,1})) \\ e_{0,1} &= \text{Enc}_{0,1}(s_{0,0}, s_{0,1}, s_{0,2}) \\ &= \text{Enc}_{0,1}(\text{Rec}_0(e_{1,0}, e_{2,0}, e_{3,0}, e_{4,0}, e_{5,0})) \end{aligned}$$

The encoded symbol  $e_{1,0}$  is computed by encoding the stored symbols  $s_{1,0}$ ,  $s_{1,1}$ , and  $s_{1,2}$ , all of which can be reconstructed from other encoded symbols  $e_{0,1}$ ,  $e_{2,1}$ ,  $e_{3,1}$ ,  $e_{4,1}$ , and  $e_{5,1}$  based on single failure recovery. Thus,  $e_{1,0}$  can be expressed as a function of encoded symbols. The encoded symbol  $e_{0,1}$  is expressed in a similar way. Now, we have two equations with two unknowns  $e_{1,0}$  and  $e_{0,1}$ . If these two equations are linearly independent, we can solve for  $e_{1,0}$  and  $e_{0,1}$ . Then we can apply  $\text{Rec}_i$  to reconstruct the lost stored symbols of node  $N_i$ . In general, to recover  $t$  failed nodes, we have a total of  $t(t-1)$  virtual symbols. We can compose  $t(t-1)$  equations based on the above idea. If these  $t(t-1)$  equations are linearly independent, we can solve for the virtual symbols. A subtle issue is that the system of equations may be unsolvable. We explain how we generalize our baseline approach for such an issue in the next subsection.

### B. Recovering Any Failure Pattern

We seek to express the virtual symbols as a function of real symbols by solving a system of equations. However, we note that for some failure patterns (i.e., the set of failed nodes), the system of equations cannot return a unique solution. A failure pattern is said to be *good* if we can uniquely express the virtual symbols as a function of the real symbols, or *bad* otherwise. Our goal is to reduce the recovery bandwidth even for bad failure patterns.

We first evaluate the likelihood of having bad failure patterns for different choices of parameters. Given an  $(n, k)$  code and  $t$  failures, there are  $\binom{n}{t}$  possible failure patterns. We enumerate all such possible failure patterns and check if each of them is bad. In practice, each stripe has a limited number of nodes (i.e.,  $n$  will not be too large) [26], [31], so we can feasibly enumerate all possible failure patterns and identify the bad ones in advance. We conduct our enumeration for both IA and PM codes.

Figure 5 shows the proportions of bad failure patterns for different combinations of  $(n, k)$  and  $t$ . We observe that among all parameters we consider, bad failure patterns only account for a small proportion, with at most 0.9% and 1.6% for IA and PM codes, respectively. Also, for some sets of parameters, we do not find any bad failure patterns. Nevertheless, we would like to reduce the recovery bandwidth for such bad failure patterns even though they are rare.

We now extend our baseline approach of CORE to deal with the bad failure patterns, with an objective of reducing

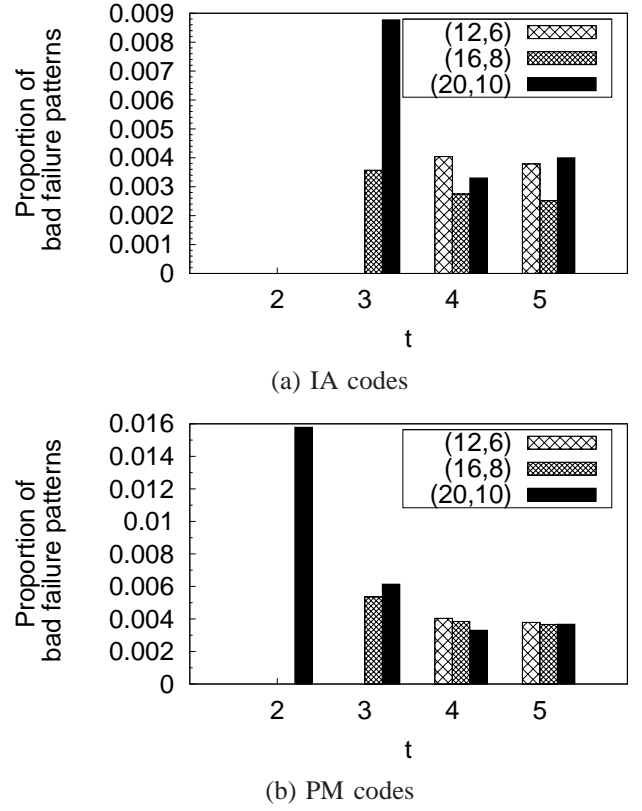


Fig. 5. Proportions of bad failure patterns for different  $(n, k)$  and  $t$ .




Fig. 6. An example of using a virtual failure pattern for a (6,3) code. If the original failure pattern  $\{N_0, N_1\}$  is bad, then we can instead recover the virtual failure pattern  $\{N_0, N_1, N_2\}$  and only download encoded symbols from nodes  $N_3, N_4, N_5$ .

the recovery bandwidth over the conventional recovery approach. For a bad failure pattern  $\mathcal{F}$ , we include one additional surviving node and form a *virtual failure pattern*  $\mathcal{F}'$ , such that  $\mathcal{F} \subset \mathcal{F}'$  and  $|\mathcal{F}'| = |\mathcal{F}| + 1 = t + 1$ . Then the relay downloads the data from the  $n - t - 1$  nodes outside  $\mathcal{F}'$  needed for reconstructing the lost data of  $\mathcal{F}'$ , although actually only the lost data of  $\mathcal{F}$  needs to be reconstructed. Figure 6 shows an example of how we use a virtual failure pattern for recovery. If  $\mathcal{F}'$  is still a bad failure pattern, then we include an additional surviving node into  $\mathcal{F}'$ , and repeat until a good failure pattern is found. Note that the size of  $\mathcal{F}'$  must be upper-bounded by  $n - k$ , as we can always connect to  $k$  surviving nodes to reconstruct the original data due to the MDS code property.

### C. Theoretical Results

In this subsection, we present two theorems. The first one shows the lower bound of recovery bandwidth. The second one shows that CORE achieves the lower bound for good failure patterns. The proofs are in Appendix.

*Theorem 1:* Suppose that we recover  $t$  failed nodes. The lower bound of recovery bandwidth is:

$$\begin{cases} \frac{Mt(n-t)}{k(n-k)} & \text{where } t < k, \\ M & \text{where } t \geq k. \end{cases} \quad \square$$

*Theorem 2:* CORE, which builds on MSR codes for single failure recovery, achieves the lower bound in Theorem 1 if we recover a good failure pattern.  $\square$


Since most failure patterns are good (with at least 99.1% and 98.4% for IA and PM codes, respectively), we conclude that CORE minimizes recovery bandwidth for a majority of failure patterns. In the next subsection, we show the actual bandwidth saving of CORE in both good and failure patterns.

### D. Analysis of Bandwidth Saving


We now study the bandwidth saving of CORE over conventional recovery. We compute the bandwidth ratio, defined as the ratio of recovery bandwidth of CORE to that of conventional recovery. We vary  $(n, k)$  and the number  $t$  of failed nodes to be recovered.

We first consider good failure patterns. For CORE, the recovery bandwidth achieves the lower bound derived in Theorem 1, and we can directly apply the theoretical results. For conventional recovery, the recovery bandwidth is the amount of original data being stored. Figure 7(a) shows the bandwidth ratio. We observe that CORE achieves bandwidth saving in both single and concurrent failures. For single failures (i.e.,  $t = 1$ ), CORE directly benefits from existing regenerating codes, and saves the recovery bandwidth by 70-80%. For concurrent failures (i.e.,  $t > 1$ ), CORE also shows the bandwidth saving, for example by 44-64%, 25-49%, and 11-36% for  $t = 2$ ,  $t = 3$  and  $t = 4$ , respectively. The bandwidth saving decreases as  $t$  increases, since more lost data needs to be reconstructed and we need to retrieve nearly the amount of original data stored. On the other hand, the bandwidth saving increases with the values of  $(n, k)$ . For example, the saving is 36-64% in (20,10) when  $2 \leq t \leq 4$ .

We now study how CORE performs for bad failure patterns. Recall from Section IV-B for each bad failure pattern  $\mathcal{F}$ , CORE forms a virtual failure pattern  $\mathcal{F}'$  that is a good failure pattern. We compute the recovery bandwidth for  $\mathcal{F}'$  based on our theoretical results in Section IV-C. Figure 7(b) shows the bandwidth ratio. We find that in all cases we consider, it suffices to add one surviving node into  $\mathcal{F}'$  (i.e.,  $|\mathcal{F}'| = |\mathcal{F}| + 1$ ) and obtain a good failure pattern. Thus, the recovery bandwidth of CORE for a bad  $t$ -failure pattern is always equivalent to that for a good  $(t + 1)$ -failure pattern. From the figure, we still see bandwidth saving of CORE over conventional recovery. For example, the saving is 25-49% in (20,10) when  $2 \leq t \leq 4$ .



(a) Good failure patterns



(b) Bad failure patterns

Fig. 7. Ratio of recovery bandwidth of CORE to that of conventional recovery.

## V. IMPLEMENTATION

We complement our theoretical analysis with prototype implementation. As a proof of concept, we implement CORE as an extension to the Hadoop Distributed File System (HDFS) [40]. We modify the source code of HDFS and its erasure code module HDFS-RAID [18].

### A. Overview of HDFS-RAID

By default, HDFS uses 3-way replication to achieve data availability. To provide data availability with smaller storage overhead, HDFS-RAID is designed to convert replicas into erasure-coded data and stripe the erasure-coded data across different nodes. We call it the *striping* operation.

HDFS-RAID uses a distributed RAID file system (DRFS) that manages the erasure-coded data stored in HDFS. In the original HDFS design, the basic data unit of the read/write operation is called a block (see Section II-A). There are a single NameNode and multiple DataNodes. The NameNode stores the metadata for HDFS blocks, while the DataNodes store HDFS blocks. On top of HDFS, HDFS-RAID adds a new node called the RaidNode, which performs the striping operation. It also periodically checks any lost blocks, and if needed, performs the recovery operation for those blocks. Also, HDFS-RAID provides a client-side interface called DRFS client, which handles all read/write requests for the erasure-coded data stored in HDFS. If a lost block is requested, then it performs degraded reads to the lost block. Both the RaidNode and the DRFS client have an ErasureCode module, which



performs the encoding/decoding operations for the erasure-coded data.

The striping operation is carried out as follows. For a given  $(n, k)$ , the RaidNode first downloads a group of  $k$  blocks (from one of the replicas for each block). It then encodes the  $k$  blocks into  $n$  blocks on a per-stripe basis (see Section II-A). The  $n$  blocks are then placed on  $n$  DataNodes. Unused replicas of the  $k$  blocks will later be removed from HDFS. The RaidNode repeats the same process for another group of  $k$  blocks.

### B. Integration into HDFS-RAID

To integrate our relay model into HDFS-RAID, we can simply deploy a relay daemon in the RaidNode and the DRFS client for failure recovery and degraded reads, respectively. CORE is implemented on HDFS release 0.22.0 with HDFS-RAID enabled. We modify both the RaidNode and the DRFS client accordingly to support concurrent recovery. Since regenerating codes need DataNodes to generate encoded symbols during recovery, we add a signal handler in each DataNode to respond to the request of encoded symbols. During recovery, the RaidNode or the DRFS client tell the surviving DataNodes about the identities of the failed nodes, and the DataNodes accordingly generate the encoded symbols.

**Optimizations of coding.** In our current prototype, we implement RS codes [34] and IA codes [42] as candidates of erasure codes and regenerating codes, respectively. We implement them in the ErasureCode module of HDFS-RAID. To minimize the computational overhead of the encoding/decoding operations, we implement the coding schemes in C++ using the Jerasure library [31], and have the ErasureCode module execute a specific coding scheme through the Java Native Interface (note that HDFS-RAID is written in Java). For each code we implement, we add *XOR transformation* [3], which changes all encoding/decoding operations into purely XOR operations, and *XOR scheduling* [17], which reduces the number of redundant XOR operations during encoding/decoding. Both XOR transformation and XOR scheduling are available in the Jerasure library [31].

**Pipelined model.** The original HDFS-RAID uses a single-threaded implementation. For further speedup, we implement a *pipelined* model that leverages multi-threading to parallelize the encoding/decoding operations. Figure 8 shows the implementation of our pipelined design in CORE, assuming that a single failure is to be recovered. The RaidNode requests meta-data from the NameNode (Steps 1-2) and downloads blocks from the surviving nodes (Steps 3-4). Then the RaidNode reconstructs the lost data using the pipelined implementation, which is composed of three stages. First, we have an *input thread* that collects data from the surviving DataNodes. The input thread then dispatches the data via a shared ring buffer to the *worker thread*, which reconstructs the lost data for the failed nodes. In the case of regenerating codes, the worker thread fetches the encoded symbols of one stripe from the ring buffer. It decodes the encoded symbols corresponding to the stripe and reconstructs the lost strips for the failed nodes. It sends the reconstructed strips to an *output thread*, and

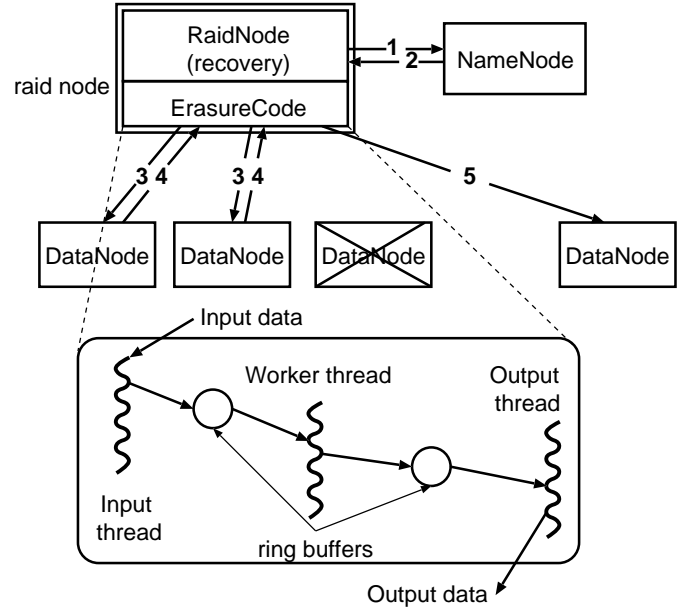


Fig. 8. Illustration of the pipelined implementation in CORE for the recovery operation, assuming that we recover a single failure. The same implementation applies to striping (in the RaidNode) and degraded reads (in the DRFS client).

processes another stripe. The output thread then collects all reconstructed stripes and uploads the resulting blocks (Step 5).

## VI. PROTOTYPE EXPERIMENTS

We experiment CORE on a distributed storage system testbed. A major deployment issue is that the overall recovery performance is determined by a combination of factors including network bandwidth, disk I/Os, encoding/decoding overhead. We address the following questions:

- Does minimizing recovery bandwidth play a key role in improving the overall recovery performance (see Section VI-A)?
- Can CORE preserve the performance of the normal striping operation offered by HDFS-RAID (see Section VI-B)?
- How much can CORE improve the recovery and degraded read performance (see Sections VI-C and VI-D)?

We conduct our experiments on an HDFS testbed with one NameNode and up to 20 DataNodes being used. Each node runs on a quad-core PC equipped with an Intel Core i5-2400 3.10GHz CPU, 8GB RAM, and a Seagate ST31000524AS 7200RPM 1TB SATA harddisk. All machines are equipped with a 1Gb/s Ethernet card and interconnected over a 1Gb/s Ethernet switch. They all run Linux Ubuntu 12.04.

We compare RS codes [34], which use conventional recovery, and CORE, which builds on IA codes [42] (see Section V-B). Both codes are implemented in C++ and compiled with GCC 4.6.3 with the -O3 option.

Our microbenchmark results (see Section VI-A) are averaged over 10 runs, while other macrobenchmark results (see Sections VI-B, VI-C, and VI-D) are averaged over five runs.



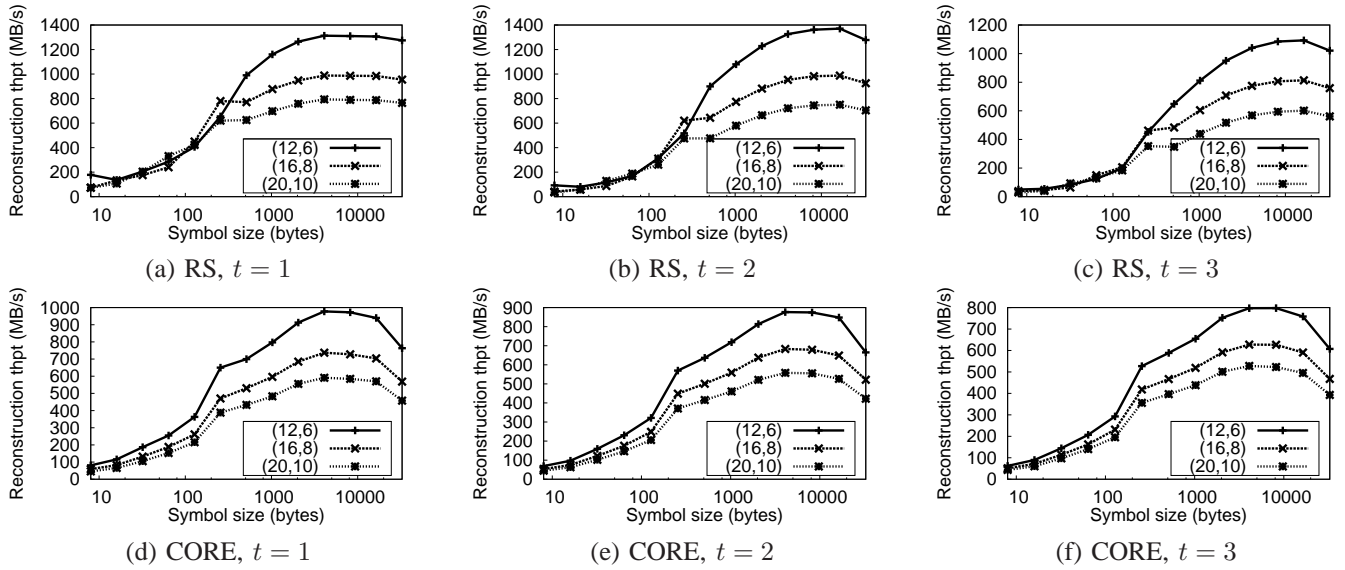


Fig. 9. Reconstruction throughput of RS codes and CORE versus the symbol size for different  $(n, k)$ .

#### A. Microbenchmark Studies

In this subsection, we conduct microbenchmark studies on the recovery operation. We first evaluate the encoding/decoding performance versus the symbol size. We then provide a breakdown analysis on different recovery steps.

**Encoding/decoding performance in reconstruction.** To evaluate the computational encoding/decoding overhead of RS codes and CORE in recovery, we measure how fast the relay decodes the symbols downloaded from surviving nodes and reconstructs the lost data. Since the encoding/decoding operations are performed over symbols (see Section II-A), our goal here is to study how the symbol size affects the encoding/decoding performance in reconstruction.

We vary the symbol size from 8 bytes to 32KB. Our evaluation operates on 30 stripes of data for different sets of  $(n, k)$ . To stress test the computational encoding/decoding performance, we eliminate the impact of disk I/Os by first loading the data that is to be downloaded by the relay for recovery into memory. We then measure the time for performing all encoding/decoding operations on the in-memory data for reconstruction. We compute the *reconstruction throughput*, which is defined as the size of the lost data divided by the reconstruction time.

Figure 9 shows the reconstruction throughput for one to three failures for RS codes and CORE. Larger  $(n, k)$  implies more failures can be tolerated, but has smaller reconstruction throughput since the generator matrix becomes larger and there is higher encoding/decoding overhead. Note that the throughput trend versus the symbol size also conforms to the results of different erasure codes in the study [31]. The throughput initially increases with the symbol size, and reaches maximum when the symbol size is around 4KB to 8KB. When the symbol size further increases, the throughput drops because of cache misses [31].

RS codes have higher reconstruction throughput than CORE

(which builds on IA codes). The reason is that the strip size of regenerating codes is  $r = n - k$  (see Section II-C), while we can implement erasure codes with  $r = 1$ . For the same  $(n, k)$ , the generator matrix of regenerating codes is larger than that of erasure codes (see Section II-A). Nevertheless, in all cases we consider, CORE has at least 500MB/s of reconstruction throughput at symbol size 8KB. Our following benchmark results show that the reconstruction performance is *not* the bottleneck in the recovery operation.

**Breakdown analysis.** Recall from Figure 2 that a recovery operation can be decomposed into five different steps. We now conduct a simplified analysis on the expected performance of each recovery step in RS codes and CORE. Our goal is to identify the bottleneck, and hence justify the need of minimizing recovery bandwidth.

We fix the storage capacity of each node to be 1GB. Suppose that we recover  $t$  failed nodes with a total of  $t$ GB of data, and that  $(n, k) = (20, 10)$  is used. We collect the system parameters based on the measurements on our testbed hardware, and derive the expected time for each recovery step as shown in Table II. We elaborate our derivations as follows.

- *I/O step.* In both RS codes and CORE, each surviving node reads all its stored data. For our disk model, our measurements (using the Linux command `hdparm`) indicate that the disk read speed is 116MB/s. Suppose that all surviving nodes read data in parallel. In the I/O step, both schemes take  $1\text{GB} \div 116\text{MB/s} \approx 8.83\text{s}$ .
- *Encode step.* In RS codes, surviving nodes do not perform encoding, while in CORE, surviving nodes encode their stored data. Suppose that all surviving nodes perform the encode step in parallel. Our measurements indicate that the encoding time on an i5-2400 machine is no more than 0.4 seconds for 1GB of raw data.
- *Download step.* The relay downloads data from other surviving nodes via its 1Gb/s interface, so its effec-

TABLE II

TIME COMPARISONS FOR DIFFERENT RECOVERY STEPS IN RS CODES AND CORE IN (20,10), ASSUMING 1GB DATA PER NODE.

time(s)	RS $t = 1$	RS $t = 2$	RS $t = 3$	CORE $t = 1$	CORE $t = 2$	CORE $t = 3$
I/O	8.83	8.83	8.83	8.83	8.83	8.83
Encode	0	0	0	0.12	0.23	0.35
Download	81.92	81.92	81.92	15.56	29.49	41.78
Reconstruct	1.30	2.75	5.17	1.75	3.69	5.87
Upload	8.19	16.38	24.58	8.19	16.38	24.58

tive transfer rate must be upper bounded by 1Gb/s (or 125MB/s). For RS codes, the relay always downloads the same amount of original data, which is  $k \times 1\text{GB} = 10\text{GB}$ . For CORE, we consider only the good failure patterns, which account for the majority of cases (see Section IV-D). From Theorem 1, the relay downloads  $0.1t(20 - t)\text{GB}$  of data (where  $t < k = 10$ ). We can derive the (minimum) download times for RS codes and CORE accordingly. In reality, the effective transfer rate is lower than 1Gb/s and the download times will be higher.

- *Reconstruction step.* We fix the symbol size at 8KB, in which both RS codes and CORE can achieve high reconstruction throughput according to our previous experiments. The reconstruction throughput values of RS codes are 594-789MB/s, while those of CORE are 523-585MB/s. We derive the reconstruction times by dividing  $t\text{GB}$  by the reconstruction throughput for  $t$  failures.
- *Upload step.* The relay uploads  $t\text{GB}$  of reconstructed data via its 1Gb/s interface. We derive the upload times as in the download step.

From our derivations, we see that the download step uses the most time among all operations. Since we can pipeline the download, reconstruction, and upload steps in the relay, we can see that the download step is the bottleneck. This justifies the need of minimizing recovery bandwidth, which we define as the amount of data transferred in the download step.

### B. Striping

We now evaluate the striping operation that is originally provided by HDFS-RAID when encoding replicas with RS codes and IA codes (used by CORE). We also compare our pipelined implementation with the original single-threaded implementation in HDFS-RAID. Our goal is to show that CORE, when using IA codes, maintains the striping performance when compared to RS codes.

For a given  $(n, k)$ , we configure our HDFS testbed with  $n$  DataNodes, one of which also deploys the RaidNode. We prepare a  $k\text{GB}$  of original data as our input. By our observation, the input size is large enough to give a steady throughput. HDFS first stores the file with the default 3-replication scheme. Then the RaidNode stripes the replica data into encoded data using either RS codes or IA codes. The encoded data is stored in  $n$  DataNodes. We rotate node identities when we place the blocks so that the parity blocks are evenly distributed across different DataNodes to achieve load balancing. We fix the symbol size at 8KB. We use the




Fig. 10. Striping throughput.

default HDFS block size at 64MB, but for some  $(n, k)$ , we alter the block size slightly to make it a multiple of the strip size (which is  $(n - k) \times 8\text{KB}$ ) for IA codes. We measure the *striping throughput* as the original size of data divided by the total time for the entire striping operation.

Figure 10 shows the striping throughput results. By parallelizing the data transfer and encoding/decoding steps, our pipelined implementation improves the striping throughput by around 50% over the original single-threaded implementation in HDFS-RAID. We see that IA codes have smaller striping throughput than RS codes in both implementations. In single-threaded implementation, IA codes have higher encoding/decoding overhead and hence show worse performance. In pipelined implementation, IA codes have strip size  $r = n - k$  and contain more symbols per stripe than RS codes with strip size  $r = 1$ . Our pipelined implementation will not start the encoding thread until the RaidNode downloads the first stripe of symbols for each group of  $k$  blocks (see Section V-A). Thus, RS codes benefit more from parallelism. However, the throughput drop in IA codes is small, by at most 6.1% only in our pipelined implementation.

### C. Recovery

We evaluate the recovery performance. We first stripe encoded data across DataNodes as in Section VI-B. Then we manually delete all blocks stored on  $t$  DataNodes to mimic  $t$  failures, where  $t = 1, 2, 3$ . Since we rotate node identities when we stripe data, the lost blocks of the  $t$  failed DataNodes include both data and parity blocks. The RaidNode recovers the failures and uploads reconstructed blocks to new DataNodes (same as the failed DataNodes in our evaluation). Here, we deploy the RaidNode in one of the new DataNodes. We measure the *recovery throughput* as the total size of lost blocks divided by the total recovery time.

Figure 11 shows the recovery throughput results. Both CORE and RS codes see higher throughput for larger  $t$  as more lost blocks are recovered. Overall, CORE shows significantly higher throughput than RS codes. The throughput gain is the highest in (20,10). For example, for single failures, the gain is  $3.45\times$ ; for concurrent failures, the gains are  $2.33\times$  and  $1.75\times$  for  $t = 2$  and  $t = 3$ , respectively.




Fig. 11. Recovery throughput.

Our experimental results are fairly consistent with our analytical results in Section IV-D. For example, in (20,10), the ratio of the reconstruction bandwidth of CORE to that of erasure codes for  $t = 2$  and  $t = 3$  are 0.36 and 0.51, respectively (see Figure 7(a)). These results translate to the recovery throughput gains of CORE at  $2.78\times$  and  $1.96\times$ , respectively. Our experimental results show slightly less gains, mainly due to disk I/O and encoding/decoding overheads that are not captured in the recovery bandwidth.

#### D. Degraded Reads

We further evaluate the degraded read performance in the presence of transient failures. The evaluation setting is the same as that of the recovery operation described in Section VI-C, except that the degraded read operation is now performed by the DRFS client. Suppose that  $t$  nodes fail, where  $t = 1, 2, 3$ . We have the DRFS client request a lost HDFS block on one of the failed DataNodes. The lost block will be reconstructed from the data of other surviving DataNodes. Here, we deploy the DRFS client in one of the failed DataNodes. We measure the *degraded read throughput*, defined as the amount of data being requested divided by the response time.

Figure 12 shows the degraded read throughput results. RS codes keep almost the same throughput for each  $(n, k)$ , as they always download  $k$  blocks for reconstruction. Overall, CORE shows a throughput gain in degraded reads. For example, if we consider the (20,10) code, CORE shows degraded throughput gain of  $3.75\times$ ,  $2.34\times$  and  $1.70\times$  for  $t = 1$ ,  $t = 2$ , and  $t = 3$ , respectively.

We point out that our concurrent reconstruction is optimized for reconstructing  $t$  lost blocks on  $t$  failures. If only one lost block is reconstructed while  $t > 1$ , it is possible to use even less reconstruction bandwidth. Nevertheless, our results still show the improvements of our concurrent reconstruction over the conventional one.

## VII. RELATED WORK

We review related work on the recovery problem for erasure codes and regenerating codes.




Fig. 12. Degraded read throughput.

**Efficient recovery.** There have been extensive studies on improving the recovery performance of coded storage systems. Muntz and Lui [28] use parity declustering (i.e., distributing data over a larger number of disks) to reduce performance degradation during recovery. Holland *et al.* [19] analyze the impact of parity declustering, and further propose workflow parallelization to speed up reconstruction. FARM [48] employs declustering to improve both recovery and reliability in large-scale storage systems. Total Recall [2] proposes a lazy repair scheme to reduce the data transferred when recovering concurrent failures. Some studies [41], [43], [45] leverage workload characteristics or access patterns to improve the reconstruction performance. Note that all the above approaches build on our definition of conventional recovery, in which first reconstruct the original data. They retrieve the size of original data, and incur high I/Os and bandwidth in general.

**Minimizing I/Os.** Several studies focus on minimizing I/Os required for recovering a single failure in erasure codes. Their approaches mainly focus on a disk array system where the disk access is the bottleneck. Authors of [44], [47], [49] propose optimal single failure recovery for RAID-6 codes. Khan *et al.* [26] show that finding the optimal recovery solution for arbitrary erasure codes is NP-hard, and propose an enumeration-based recovery algorithm. They also propose a modified Reed-Solomon code for efficient degraded reads. Authors of [50], [51] propose greedy heuristics to speed up the search of solutions for single failure recovery. Note that the performance gains of the above solutions over the conventional recovery are generally less than 30%, while regenerating codes achieve a much higher gain in single failure recovery (see Section VI).

Huang *et al.* [22] propose local recovery codes that reduce the bandwidth and I/O when reconstructing a lost data fragment. They evaluate the codes atop the Windows Azure Storage system. Sathiamoorthy *et al.* [35] also propose local recovery codes, and evaluate the codes atop HDFS-RAID as in our work. It is worth noting that the codes in both studies [22], [35] are non-MDS codes with additional parities added to storage, so as to trade for better recovery performance. Both studies focus on optimizing single failure recovery. Our work differs from them in several aspects: (i) we consider optimal



minimum storage regenerating codes that are MDS codes, (ii) we consider recovering both single and concurrent failures, (iii) we experiment regenerating codes that require storage nodes to perform encoding operations.

**Minimizing recovery bandwidth.** Regenerating codes [11] minimize the recovery bandwidth for a single failure in a distributed storage system. In contrast with the above solutions that minimize I/Os, most regenerating codes typically read all stored data to generate encoded data (e.g., [7], [11], [32], [37], [42], [46]). Some regenerating codes do not send encoded data during recovery (e.g., [30], [33]), but generally have higher storage overhead.

**Cooperative recovery.** Several theoretical studies (e.g., [21], [25], [38], [39]) address concurrent failure recovery based on regenerating codes, and they focus on recovery of lost data on new nodes. They all consider a *cooperative model*, in which the new nodes exchange among themselves their data being read from surviving nodes during recovery. Authors of [21], [25] prove that the cooperative model achieves the same optimal recovery bandwidth as ours, but they do not provide explicit constructions of regenerating codes that achieve the optimal point. Authors of [38], [39] provide such explicit implementations, but they focus on limited parameters and the resulting implementations do not provide any bandwidth saving over erasure codes. A drawback of the cooperative model requires coordination among the new nodes to perform recovery, and its implementation complexities are unknown. Extending it for degraded reads is also non-trivial, as clients simply request lost data instead of recovering lost data on new nodes.

**Implementation of regenerating codes.** Implementation studies of regenerating codes recently receive attention from the research community, such as [12], [20], [23], [24]. In particular, Jiekak *et al.* [24] evaluate the encoding/decoding performance of PM codes. Note that the studies [12], [23], [24] do not integrate regeneration codes into a real storage system. NCCloud [20] is a multiple-cloud storage prototype based on regenerating codes, but it only implements non-systematic regenerating codes. We point out that existing implementation studies only focus on single failure recovery.

## VIII. CONCLUSIONS AND FUTURE WORK

We address the reconstruction problem in a distributed storage system in the presence of single and concurrent failures, from both theoretical and applied perspectives. We explore the use of regenerating codes (or network coding) to provide fault-tolerant storage and minimize the bandwidth of data transfer during reconstruction. We propose a system CORE, which generalizes existing optimal single-failure-based regenerating codes to support the recoveries of both single and concurrent failures. We theoretically show that CORE minimizes the reconstruction bandwidth in most concurrent failure patterns. Our scheme adopts a relay model that can be easily integrated into real storage systems. To demonstrate, we prototype CORE as a layer atop Hadoop HDFS, and show

via testbed experiments that we can speed up both recovery and degraded read operations.

**Future work.** Since we currently implement CORE on HDFS, one interesting issue is to analyze how CORE affects the performance of MapReduce [9] jobs in the presence of failures. Also, we plan to explore the implementation of CORE in wide-area storage systems (e.g., [2], [6], [8], [27]) in future work.

**Availability.** We plan to release the source code of our CORE prototype when the final version of the paper is published.

## REFERENCES

- [1] R. Ahlswede, N. Cai, S. Li, and R. Yeung. Network Information Flow. *IEEE Trans. on Information Theory*, 46(4):1204–1216, Jul 2000.
- [2] R. Bhagwan, K. Tati, Y. Cheng, S. Savage, and G. Voelker. Total Recall: System Support for Automated Availability Management. In *Proc. of USENIX NSDI*, Oct 2004.
- [3] J. Bloemer, M. Kalfane, R. Karp, M. Karpinski, M. Luby, and D. Zuckerman. An XOR-based Erasure-Resilient Coding Scheme. Technical report, The International Computer Science Institute, Berkeley, CA., Aug 1995.
- [4] B. Calder, J. Wang, A. Ogus, N. Nilakantan, A. Skjolsvold, S. McKelvie, Y. Xu, S. Srivastav, J. Wu, H. Simitci, et al. Windows Azure Storage: A Highly Available Cloud Storage Service with Strong Consistency. In *Proc. of ACM SOSP*, Oct 2011.
- [5] B. Chun, F. Dabek, A. Haeberlen, E. Sit, H. Weatherspoon, M. Kaashoek, J. Kubiatowicz, and R. Morris. Efficient Replica Maintenance for Distributed Storage Systems. In *Proc. of USENIX NSDI*, May 2006.
- [6] Cleversafe. <http://www.cleversafe.com>.
- [7] D. Cullina, A. Dimakis, and T. Ho. Searching for Minimum Storage Regenerating Codes. In *Proc. of Allerton Conf.*, Sep 2009.
- [8] F. Dabek, M. Kaashoek, D. Karger, R. Morris, and I. Stoica. Wide-Area Cooperative Storage with CFS. *ACM SIGOPS Operating Systems Review*, 35(5):202–215, Dec 2001.
- [9] J. Dean and S. Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. In *Proc. of USENIX OSDI*, Dec 2004.
- [10] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Voshall, and W. Vogels. Dynamo: Amazon’s Highly Available Key-Value Store. In *Proc. of ACM SOSP*, 2007.
- [11] A. Dimakis, P. Godfrey, Y. Wu, M. Wainwright, and K. Ramchandran. Network Coding for Distributed Storage Systems. *IEEE Trans. on Information Theory*, 56(9):4539–4551, Sep 2010.
- [12] A. Duminuco and E. Biersack. A Practical Study of Regenerating Codes for Peer-to-Peer Backup Systems. In *Proc. of IEEE ICDCS*. IEEE, Jun 2009.
- [13] D. Ford, F. Labelle, F. I. Popovici, M. Stokel, V.-A. Truong, L. Barroso, C. Grimes, and S. Quinlan. Availability in Globally Distributed Storage Systems. In *Proc. of USENIX OSDI*, Oct 2010.
- [14] S. Ghemawat, H. Gobioff, and S. Leung. The Google File System. In *Proc. of ACM SOSP*, Dec 2003.
- [15] K. M. Greenan, E. L. Miller, and T. J. E. Schwarz. Optimizing Galois Field Arithmetic for Diverse Processor Architectures and Applications. In *Proc. of IEEE MASCOTS*, 2008.
- [16] A. Haeberlen, A. Mislove, and P. Druschel. Glacier: Highly Durable, Decentralized Storage Despite Massive Correlated Failures. In *Proc. of USENIX NSDI*, May 2005.
- [17] J. Hafner, V. Deenadhayalan, K. Rao, and J. Tomlin. Matrix Methods for Lost Data Reconstruction in Erasure Codes. In *Proc. of USENIX FAST*, Dec 2005.
- [18] HDFS-RAID. <http://wiki.apache.org/hadoop/HDFS-RAID>.
- [19] M. Holland, G. Gibson, and D. Siewiorek. Architectures and Algorithms for On-Line Failure Recovery in Redundant Disk Arrays. *Distributed and Parallel Databases*, 2(3):295–335, Jul 1994.
- [20] Y. Hu, H. Chen, P. Lee, and Y. Tang. NCCloud: Applying Network Coding for the Storage Repair in a Cloud-of-Clouds. In *Proc. of USENIX FAST*, Feb 2012.

- [21] Y. Hu, Y. Xu, X. Wang, C. Zhan, and P. Li. Cooperative Recovery of Distributed Storage Systems from Multiple Losses with Network Coding. *IEEE Journal on Selected Areas in Communications (JSAC)*, 28(2):268–276, Feb 2010.
- [22] C. Huang, H. Simitci, Y. Xu, A. Ogus, B. Calder, P. Gopalan, J. Li, and S. Yekhanin. Erasure Coding in Windows Azure Storage. In *Proc. of USENIX ATC*, Jun 2012.
- [23] Z. Huang, E. Biersack, and Y. Peng. Reducing Repair Traffic in P2P Backup Systems: Exact Regenerating Codes on Hierarchical Codes. *ACM Trans. on Storage*, 7(3):10, Oct 2011.
- [24] S. Jieka, A.-M. Kermarrec, N. L. Scouarnec, G. Straub, and A. Van Kempen. Regenerating Codes: A System Perspective. *CoRR*, abs/1204.5028, 2012.
- [25] A. Kermarrec, N. Le Scouarnec, and G. Straub. Repairing Multiple Failures with Coordinated and Adaptive Regenerating Codes. In *Proc. of NetCod*, Jun 2011.
- [26] O. Khan, R. Burns, J. Plank, W. Pierce, and C. Huang. Rethinking Erasure Codes for Cloud File Systems: Minimizing I/O for Recovery and Degraded Reads. In *Proc. of USENIX FAST*, Feb 2012.
- [27] J. Kubiatiowicz, D. Bindel, Y. Chen, S. Czerwinski, P. Eaton, D. Geels, R. Gummadi, S. Rhea, H. Weatherspoon, W. Weimer, C. Wells, and B. Zhao. OceanStore: An Architecture for Global-Scale Persistent Storage. In *Proc. of ACM ASPLOS-IX*, Nov 2000.
- [28] R. Muntz and J. Lui. Performance Analysis of Disk Arrays under Failure. In *Proceedings of the 16th International Conference on Very Large Data Bases*, Aug 1990.
- [29] S. Nath, H. Yu, P. B. Gibbons, and S. Seshan. Subtleties in Tolerating Correlated Failures in Wide-area Storage Systems. In *Proc. of USENIX NSDI*, May 2006.
- [30] D. Papailiopoulos, J. Luo, A. Dimakis, C. Huang, and J. Li. Simple Regenerating Codes: Network Coding for Cloud Storage. In *Proc. of IEEE INFOCOM*, Mar 2012.
- [31] J. Plank, J. Luo, C. Schuman, L. Xu, and Z. Wilcox-O’Hearn. A Performance Evaluation and Examination of Open-Source Erasure Coding Libraries for Storage. In *Proc. of USENIX FAST*, Feb 2009.
- [32] K. Rashmi, N. Shah, and P. Kumar. Optimal Exact-Regenerating Codes for Distributed Storage at the MSR and MBR Points via a Product-Matrix Construction. *IEEE Trans. on Information Theory*, 57(8):5227–5239, Aug 2011.
- [33] K. Rashmi, N. Shah, P. Kumar, and K. Ramchandran. Explicit Construction of Optimal Exact Regenerating Codes for Distributed Storage. In *Proc. of Allerton Conf.*, Sep 2009.
- [34] I. Reed and G. Solomon. Polynomial Codes over Certain Finite Fields. *Journal of the Society for Industrial and Applied Mathematics*, 8(2):300–304, Jun 1960.
- [35] M. Sathiamoorthy, M. Asteris, D. Papailiopoulos, A. G. Dimakis, R. Vadali, S. Chen, and D. Borthakur. XORing Elephants: Novel Erasure Codes for Big Data. *Proceedings of the VLDB Endowment (to appear)*, 2013.
- [36] B. Schroeder and G. A. Gibson. Disk Failures in the Real World: What Does an MTTF of 1,000,000 Hours Mean to You? In *Proc. of USENIX FAST*, Feb 2007.
- [37] N. Shah, K. Rashmi, P. Kumar, and K. Ramchandran. Interference Alignment in Regenerating Codes for Distributed Storage: Necessity and Code Constructions. *IEEE Trans. on Information Theory*, 58(99):2134–2158, Apr 2012.
- [38] K. Shum. Cooperative Regenerating Codes for Distributed Storage Systems. In *Proc. of IEEE Int. Conf. on Communications (ICC)*, Jun 2011.
- [39] K. Shum and Y. Hu. Exact Minimum-Repair-Bandwidth Cooperative Regenerating Codes for Distributed Storage Systems. In *Proc. of IEEE Int. Symp. on Information Theory (ISIT)*, Jul 2011.
- [40] K. Shvachko, H. Kuang, S. Radia, and R. Chansler. The Hadoop Distributed File System. In *Proc. of IEEE MSST*, May 2010.
- [41] M. Sivathanu, V. Prabhakaran, A. Arpaci-Dusseau, and R. Arpaci-Dusseau. Improving Storage System Availability with D-GRAID. *ACM Transactions on Storage (TOS)*, 1(2):133–170, May 2005.
- [42] C. Suh and K. Ramchandran. Exact-Repair MDS Code Construction using Interference Alignment. *IEEE Trans. on Information Theory*, 57(3):1425–1442, Mar 2011.
- [43] L. Tian, D. Feng, H. Jiang, K. Zhou, L. Zeng, J. Chen, Z. Wang, and Z. Song. PRO: A Popularity-based Multi-threaded Reconstruction Optimization for RAID-Structured Storage Systems. In *Proc. of USENIX FAST*, Feb 2007.
- [44] Z. Wang, A. Dimakis, and J. Bruck. Rebuilding for Array Codes in Distributed Storage Systems. In *IEEE GLOBECOM Workshops*, Dec 2010.
- [45] S. Wu, H. Jiang, D. Feng, L. Tian, and B. Mao. WorkOut: I/O Workload Outsourcing for Boosting RAID Reconstruction Performance. In *Proc. of USENIX FAST*, Feb 2009.
- [46] Y. Wu, A. Dimakis, and K. Ramchandran. Deterministic Regenerating Codes for Distributed Storage. In *Proc. of Allerton Conf.*, Sep 2007.
- [47] L. Xiang, Y. Xu, J. Lui, Q. Chang, Y. Pan, and R. Li. A Hybrid Approach to Failed Disk Recovery Using RAID-6 Codes: Algorithms and Performance Evaluation. *ACM Trans. on Storage*, 7(3):11, Oct 2011.
- [48] Q. Xin, E. Miller, and S. Schwarz. Evaluation of Distributed Recovery in Large-Scale Storage Systems. In *Proc. of HPDC*, Jun 2004.
- [49] S. Xu, R. Li, P. Lee, Y. Zhu, L. Xiang, Y. Xu, and J. Lui. Single Disk Failure Recovery for X-code-based Parallel Storage Systems. *IEEE Trans. on Computers*, To appear.
- [50] Y. Zhu, P. Lee, Y. Hu, L. Xiang, and Y. Xu. On the Speedup of Single-Disk Failure Recovery in XOR-Coded Storage Systems: Theory and Practice. In *Proc. of IEEE MSST*, Apr 2012.
- [51] Y. Zhu, P. Lee, L. Xiang, Y. Xu, and L. Gao. A Cost-based Heterogeneous Recovery Scheme for Distributed Storage Systems with RAID-6 Codes. In *Proc. of IEEE DSN*, Jun 2012.

## APPENDIX

### A. Proof of Theorem 1

We can formally build our proof based on the analysis of the information flow graph as in [11]. Here, we only show the main idea. Let  $d$  be the number of surviving nodes from which the relay downloads data for recovery. Let  $\beta$  be the amount of data downloaded (per stripe) from each of the  $d$  surviving nodes to recover  $t$  failed nodes. We assume that the reconstructed data will be stored on  $t$  new nodes, which contain a total of  $d\beta$  units of information.

We first consider  $t < k$ . Due to the MDS property, we can restore the original data from any  $k$  out of  $n$  nodes, each storing  $\frac{M}{k}$  units of data. For example, we can select a set of any  $k - t$  originally surviving nodes (denoted by set  $\mathcal{X}$ ) and a set of any  $t$  new nodes (denoted by set  $\mathcal{Y}$ ) for some  $t \leq t$ . The total amount of useful information must be at least  $M$  in order for the original data to be restorable. However,  $\mathcal{Y}$  contains  $(k - \hat{t})\beta$  units of information derived from  $\mathcal{X}$ . By excluding the redundant information, we require:

$$\frac{M}{k}(k - \hat{t}) + (d\beta - (k - \hat{t})\beta) \geq M, \text{ for any } \hat{t} \leq t.$$

The left side is minimum when  $\hat{t} = t$ . Thus, the recovery bandwidth (i.e.,  $d\beta$ ) must be at least  $\frac{M \times d \times t}{k(d - k + t)}$ . To minimize the recovery bandwidth with respect to  $d$ , we set  $d = n - t$  and the result follows.

When  $t \geq k$ , any  $k$  out of the  $t$  new nodes must be able to restore the original data due to the MDS property. Thus, the  $t$  new nodes must contain  $M$  units of useful information, which can be reconstructed by downloading data from any  $k$  surviving nodes as in erasure codes. The recovery bandwidth is  $M$ . ■

### B. Proof of Theorem 2

Since MSR codes achieve the lower bound of recovery bandwidth for single failure recovery, the amount of data downloaded from *each* surviving node is  $\frac{M}{k(n-k)}$  [11] (see Equation (1)).

Consider  $t < k$ . CORE in essence performs  $t$  single failure recoveries based on MSR codes, and in each recovery we actually download  $\frac{M}{k(n-k)}$  units of data from each of the  $n - t$  surviving nodes. If the failure pattern is good, then we can recover the virtual symbols and hence the lost data. The lower bound is hit for  $t < k$ . For  $t \geq k$ , we can simply download  $M$  units of data from any  $k$  surviving nodes and *any* failure pattern can be recovered. The result follows. ■